REMARKS

This Amendment is filed in response to the Office Action filed on May 2, 2007. All objections and rejections are respectfully traversed.

Claims 1- 38 are in this case.

Claims 29-38 were added to better claim the invention.

Claims 14 and 22 were amended to better claim the invention.

REQUEST FOR EXAMINER INTERVIEW

Applicant respectfully requests a telephonic interview with the Examiner after the Examiner has had an opportunity to consider this Amendment, but before the issuance of the next Office Action. The Applicant may be reached at 617-951-2500.

REJECTION UNDER 35 U.S.C. §102(e)

At paragraph 3 of the Office Action, claims 1-2, 4 and 6-12 were rejected under 35 U.S.C. §102(e) as being anticipated by Sutherland *et al.* U.S. Patent Publication No. 2002/0114341 published on August 22, 2002, (hereinafter "Sutherland"). Applicant Respectfully traverses rejection.

The present invention, as set forth in the representative claim 1 comprises:

> A storage system for use in a storage system cluster, the storage system comprising:
>
> ***a storage operating system including a cluster connection manager*** adapted to create, destroy, and maintain one or more communication sessions with a cluster partner, the cluster connection manager operatively interconnected with a set of cluster connection manager clients.

Sutherland discloses a system and method for a peer-to-peer network that includes a storage coordinator that *centrally* manages distributed resources in accordance with system policies administered through a *central* administrative console. A peer-to-peer storage network is made up of two or more storage nodes that wish to communicate with one another. The communications between the nodes are managed using a storage coordinator node. Although there may be more than one storage coordinator in a peer-to-peer storage network, the storage coordinator is itself an *additional node* with the peer-to-peer storage network (that is, each storage coordinator is a different node than each of the storage nodes in the peer-to-peer storage network). In Sutherland, each storage coordinator node is responsible for cataloging each of the storage nodes that are available in the peer-to-peer storage network and the resources available on each of the respective storage nodes. Furthermore, when a storage node wishes to access/communicate with another storage node on the peer-to-peer storage network, the requesting storage node must register/authenticate itself by first accessing the storage coordinator node. The storage coordinator node then assesses whether the communication can be made and sends its own access request to the other storage node. Therefore, all communication requests between storage nodes are sent through the storage coordinator node prior to accessing each other's storage resources.

Applicant respectfully urges that the Sutherland reference does not show Applicant's claimed novel ***a storage operating system including a cluster connection manager.***

Applicant's claimed invention is directed to a storage operation system/node which includes a cluster connection manager. The cluster connection manager is located and executes *within* storage system/node. This cluster connection manager is utilized to open, destroy and maintain communications with a cluster peer process (e.g. a failover monitor that implements various failover features). Furthermore, the cluster connection

manager is configured to establish and maintain peer-to-peer connections between its storage node and cluster partner storage node (which also has its own cluster connection manager).

Sutherland does not disclose Applicant's a storage operating system including a cluster connection manager. That is, Sutherland does not disclose using a cluster connection manager which is part of a storage operating system to create, destroy, or maintain a communication between the storage system and a cluster partner which also has its own cluster connection manager. In Sutherland, (*See* Fig.6), the storage coordinator is described as an *additional centralized node* in a peer-to-peer storage network (i.e., the storage nodes and the storage coordinator nodes are separate nodes). Applicant's invention, however, claims using a cluster connection manager included within the storage node/system to create, destroy and maintain communication sessions between the storage system and a cluster partner.(*See* Fig 3.).

Applicant respectfully urges that the Sutherland patent is legally precluded from anticipating the claimed invention under 35 U.S.C §102(e) because of the absence from Sutherland of Applicant's novel ***a storage operating system including a cluster connection manager.***

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

REJECTION UNDER 35 U.S.C. §103(a)

At paragraph 15 of the Office Action, claims 5, 13-19, 25-26 and 28 were rejected under 35 U.S.C. §103(a) as being unpatentable over Sutherland in view of Meyer *et al.* U.S. Patent No. 7,203,730 issued on April 10, 2007 (hereinafter "Meyer"). Applicant respectfully traverses this rejection.

The present invention, as set forth in the representative claim 14 comprises:

A method for initiating a peer-to-peer communication session, the method comprising the steps of:

creating, using a cluster connection manager executing on a storage system, an initial connection with a cluster partner;

exchanging a set of peer connection information;

passing a set of client information to the cluster partner;

creating a set of appropriate communication ports;

alerting the cluster partner of a ready status; and

alerting a set of clients that the cluster partner is in a ready state.

As noted above, Sutherland discloses a system and method for a peer-to-peer network that includes a storage coordinator that *centrally* manages distributed resources in accordance with system policies administered through a *central* administrative console. Although there may be more than one storage coordinator in a peer-to-peer storage network, the storage coordinator is itself an *additional node* with the peer-to-peer storage network (that is again, each storage coordinator is a different node than each of the storage nodes in the peer-to-peer storage network). Therefore the storage node and the storage coordinator node are separate entities.

Meyer discloses a SCSI device manager which manages a SCSI device. The SCSI device manager is responsible for determining the initial state of the device when it

is presented to the manager by a discovery manager. Furthermore, Meyer alerts any clients who wish to be informed when the SCSI devices come and go out of the system thereby informing the clients of the readiness status of the device.

Applicant respectfully urges that neither Sutherland nor Meyer show Applicant's claimed novel ***creating, using a cluster connection manager executing on a storage system, an initial connection with a cluster partner***;

As noted above, Applicant's claimed invention is directed to a storage operation system/node which includes a cluster connection manager. The cluster connection manager is located and executes *within* storage system/node. This cluster connection manager is utilized to open, destroy and maintain communications with a cluster peer process.

Neither Meyer, nor Sutherland disclose Applicant's a storage operating system including a cluster connection manager. Sutherland disclose discloses using a cluster connection manager which is not part of a storage operating system node. As mentioned above, Sutherland also describes the storage coordinator as an *additional centralized node* in a peer-to-peer storage network (i.e., the storage nodes and the storage coordinator nodes are separate nodes).(*See* Fig.6). Applicant's invention, however, claims using a cluster connection manager included within the storage node/system to create, destroy and maintain communication sessions between the storage system and cluster partners.(*See* Fig 3.). Meyer merely discloses alerting a client of a SCSI device manager's readiness status.

Applicant respectfully urges that Sutherland and Meyer, either taken singly or in any combination, are legally insufficient to render the presently claimed invention obvious under 35 U.S.C 103(a) because of the absence in each of the cited patents of Appli-

cant's claimed novel ***creating, using a cluster connection manager executing on a storage system, an initial connection with a cluster partner.***

CONCLUSION

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Duane H. Dreger/
Duane H. Dreger
Reg. No. 48,836
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500